


Navigator

NAVO MSRC

SPRING 2008




THE MOST INCOMPREHENSIBLE THING
ABOUT THE WORLD IS THAT IT IS AT
ALL COMPREHENSIBLE.

- ALBERT EINSTEIN

***Providing HPC Resources
to DoD Communities and Beyond...***

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 2008		2. REPORT TYPE		3. DATES COVERED 00-00-2008 to 00-00-2008	
4. TITLE AND SUBTITLE NAVO MSRC Navigator. Spring 2008				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Oceanographic Office (NAVO),Major Shared Resource Center (MSRC),1002 Balch Boulevard,Stennis Space Center,MS,39522				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 24	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			



As we approach the summer of 2008 the NAVO MSRC is flush with change and growth. For the first time in three years the Center will once again feature CRAY technology along with its powerful, reliable arsenal of IBM Power platforms. Our Technology Insertion 2008 (TI-08) enhancements—an IBM P6 and a CRAY XT5 cluster—will quadruple the amount of computing capability we provide to our users. You will find more details on these systems on page 17 of this issue of The Navigator.

Considerable infrastructure enhancements are currently being undertaken in two facilities to accommodate the CRAY XT5, the water-cooled IBM P6, and their counterpart test and development systems. Our Sun F12000 mass storage servers JULES and VINCENT will be replaced with a single Sun M5000 server within this calendar year.

We will also be bringing SL8500 technology for our Disaster Recovery (DR) capabilities to accommodate the growth of other High Performance Computing Modernization Program (HPCMP) centers' data holdings. With the award of the HPCMP's Next Generation Technical Services contract to Lockheed Martin (the incumbent contractor at NAVO MSRC) we look forward to embracing the new cooperative opportunities among our sister centers.

Although the MSRC still resides physically in the Naval Oceanographic Office (NAVOCEANO) buildings, in May 2007 the Center moved organizationally to NAVOCEANO's parent command, the Naval Meteorology and Oceanography Command (COMNAVMETOCCOM). And, for the first time

Amidst a Year of Changes, NAVO MSRC Remains Focused

Christine Cuicchi
Computational Science and Applications Lead,
NAVO MSRC

in nearly eight years, the Center is bringing aboard new government staff to enhance and broaden our expertise.

During this time of growth and transition we remain dedicated to providing reliable computational capabilities and exceptional, innovative service to our users and to the Program. We're excited about our future and your part in it.

Contents

The Naval Oceanographic Office (NAVO) Major Shared Resource Center (MSRC): Delivering Science to the Warfighter

The NAVO MSRC provides Department of Defense (DoD) scientists and engineers with high performance computing (HPC) resources, including leading edge computational systems, large-scale data storage and archiving, scientific visualization resources and training, and expertise in specific computational technology areas (CTAs). These CTAs include Computational Fluid Dynamics (CFD), Climate/Weather/Ocean Modeling and Simulation (CWO), Environmental Quality Modeling and Simulation (EQM), Computational Electromagnetic and Acoustics (CEA), and Signal/Image Processing (SIP).

NAVO MSRC
Code 0TC
1002 Balch Boulevard
Stennis Space Center, MS 39522
1-800-993-7677 or
msrchelp@navo.hpc.mil

NAVO MSRC Navigator

www.navo.hpc.mil/Navigator

NAVO MSRC Navigator is a biannual technical publication designed to inform users of the news, events, people, accomplishments, and activities of the Center. For a free subscription or to make address changes, contact NAVO MSRC at the above address.

EDITOR:

Gioia Furness Petro, petrogio@navo.hpc.mil

DESIGNERS:

Kerry Townson, ktownson@navo.hpc.mil

Lynn Yott, lynn@navo.hpc.mil

Any opinions, conclusions, or recommendations in this publication are those of the author(s) and do not necessarily reflect those of the Navy or NAVO MSRC. All brand names and product names are trademarks or registered trademarks of their respective holders. These names are for information purposes only and do not imply endorsement by the Navy or NAVO MSRC.

Approved for Public Release
Distribution Unlimited

NAVO MSRC Update

- 2 Amidst a Year of Changes, NAVO MSRC Remains Focused

Feature Articles

- 5 Hypersonic Continuum/DSMC Methodology for Steady/
Transient Missile and Re-Entry Vehicle Flowfields
- 9 Using ParaView for Remote Visualization on IBM AIX
- 12 Global and Mesoscale Ensemble Forecasts of Tropical Cyclones
- 17 EINSTEIN and DAVINCI Come to the MSRC

The Porthole

- 19 Visitors to the Naval Oceanographic Office
Major Shared Resource Center

Navigator Tools and Tips

- 21 Top 5 Frequently Asked Questions about \$HOME and
\$WORKDIR on the NAVO MSRC High Performance
Computing Systems

Upcoming Events

- 23 Coming Events

SuperComputing '07

Ingredients for Success: The Supercomputing07 HPCMP Booth

- A 40 foot by 50 foot floor space
- Over 200 feet of ethernet cabling
- Nearly forty distinct graphic displays
- 13 LCD monitors
- 19 HPCMO team members
- Countless staff across the country
- Numerous videos and on-site audio/visual support
- 15 electrical outlets
- 15 crates of booth components
- Nine NAVO team members
- An 18-foot tall user success story wall
- Nine organizations
- Dedicated website

That's what it took to prepare, power, and present the DoD HPCMP's SC07 exhibit booth, one of the largest booths on the exhibit hall floor at Supercomputing07 held in Reno, Nevada, November 10-16, 2007. The booth's installation was accomplished in record time, leaving ample time for the organizations represented to populate their respective stations with videos and distribution materials prior to the opening of the show floor. Thanks to the expanded booth layout and a high-visibility location on the SC07 exhibit floor, the HPCMP booth enjoyed a large number of visitors during the conference.

NAVO MSRC Team Members:

Randy Becnel, Lockheed Martin Enterprise Solutions and Services (LMES&S)
Sheila Carbonette, LMES&S
Bob Gahrman, Science Applications International Corporation (SAIC)
Patrick Thompson, LMES&S
John Skinner, LMES&S
Kerry Townson, SAIC
Lee Whatley, LMES&S
Lynn Yott, SAIC
Christine Cuicchi, Booth Chair, NAVO MSRC

Additional support was provided by:

Brandt Barnes, U.S. General Services Administration (GSA)
Lisa Bourg, SAIC
Chuck Breath, SAIC
Jennifer Rabert, NAVO MSRC

The DoD Team consisted of representatives from six centers and three HPCMP components, and we thank them all for their participation in the SC07 booth.

DoD Team Members:

AFRL MSRC: Joan Henley, Rachel Crouse
ARL MSRC: Denise Brown, Debbie Thompson, Brian Simmonds, Ginny To
ARSC: Debra Damon
DAAC: Paul Adams, Mike Stephens
ERDC MSRC: Scottie Swillie, Charles Ray
MHPCC: Betty Duncan, Michelle Hershey
SMDC: Jackie Steele
HPCMP: Jim Cook and Sandy Allen (DREN), Dianna LaFratte (PET/SAS), Jackie Steele, Timothy Jones

Special thanks go to Scottie Swillie and Charles Ray, both of ERDC MSRC, and Jim Cook, HPCMP DREN, for their invaluable assistance during the booth set-up and tear-down. The NAVO MSRC appreciates everyone's hard work, which guaranteed the success of the exhibition, and wishes the best of luck to Scottie Swillie as he chairs the SC08 HPCMP booth.

NAVO MSRC hosts the HPCMP booth at SC07

Hypersonic Continuum/DSMC Methodology for Steady/Transient Missile and Re-Entry Vehicle Flowfields

J.L. Papp, R.G. Wilmoth, D.B. VanGilder, and S.M. Dash, Combustion Research and Flow Technology, Inc.
K. Kennedy, U.S. Army Aviation and Missile Research, Development, and Engineering Center




Figure 1. Examples of embedded continuum/rarefied flow fields.

INTRODUCTION

Higher-altitude missile and re-entry vehicle flowfield simulations often require the unification of continuum Computational Fluid Dynamics (CFD) and Direct Simulation Monte Carlo (DSMC) methodology.

For missiles with propulsive plumes (as well as for missiles or Re-entry Vehicles (RVs) with divert/control jets), the overall flow is rarefied at altitudes above 65-75 kilometers (km), requiring the use of DSMC methods. However, the plume (or divert/control jet) is a continuum from the nozzle exit plane to a breakdown surface whose extent is based on rarefaction criterion such as that formulated by Bird.¹ Hence, a hybrid approach is required for coupling the continuum and DSMC solutions along the breakdown surface, as schematized in Figure 1(a), for a conventional rocket plume.

For basic missile/plume problems, the continuum to DSMC coupling at the

breakdown surface is one-way, with the continuum solution providing inflow boundary conditions for the DSMC solution. For more complex problems with multiple nozzles or large angles-of-attack, coupling can be two-way and thus more complicated.

For RV or aerocapture flowfield problems at moderate altitudes (50-70 km), the situation is reversed, with the flow being largely continuum with embedded rarefied flow in the afterbody/wake region that requires the use of DSMC methodology, as schematized in Figure 1(b).

For missile or RV problems with divert/control jets, breakdown surface geometries can be complex, which has led to our development of more advanced techniques for their construction in an Automated Efficient Generalized Interface Surface (AEGIS) Toolkit.

Continued Next Page...




Figure 2. Demonstration of surface generation using GDM methodology within AEGIS toolkit for RCS thruster.

Extensions to unsteady problems using hybrid methodology (such as transient jets or bursts at high altitudes) require the construction of time-varying breakdown methodology. Many problems of interest deal with plumes or divert/control jets from solid propellant motors that can contain relatively high loadings of particulates, such as Al_2O_3 .

While the continuum plume codes contain well-established methodology for the fully-coupled analysis of particulates using either Eulerian or Lagrangian numerics,² DSMC codes do not and thus have required upgrades for particulate capabilities.

The DSMC code utilized for our studies is called Plume DSMC Analysis Code (PDAC).³ PDAC is an extended version of the DSMC Analysis Code (DAC)⁴ with particulate and unsteady upgrades, as well as upgrades for plume thermo-chemistry.

AEGIS METHODOLOGY

The AEGIS Toolkit employs Geometrically Deformed Model (GDM) methodology, which implements a cost minimization algorithm based on image (breakdown parameter), topology (surface elasticity), and deformation

(forcing function) parameters, to produce a continuous and water-tight triangulated surface.

The procedure can be thought of as a balloon expanding within a cage. As the balloon is inflated (deformation), it gradually takes on the shape of the cage while the surface elasticity (topology) of the balloon prevents it from leaking out between any of the bars.

Eventually, a point is reached where the surface cannot expand without violating any of the constraints and the process is complete. So, even though the image function may be discontinuous, smoothness is maintained as a constraint, and an appropriate coupling boundary surface can be obtained from an oftentimes complex field function.

An example of the image identification and surface generation process is shown in Figure 2. The process begins by choosing an appropriate image constraint (Figure 2a) that will define the location of breakdown, which can be the Bird breakdown parameter¹ or similar variable.

The GDM process begins by introducing a seed surface within the image constraint (Figure 2b).




Figure 3. Estimated breakdown surface for multi-nozzle plume.




Figure 4. Snapshots of interface surface deformation using GDM component of AEGIS toolkit.




Figure 5. Expansion of second interface surface into far-field region.

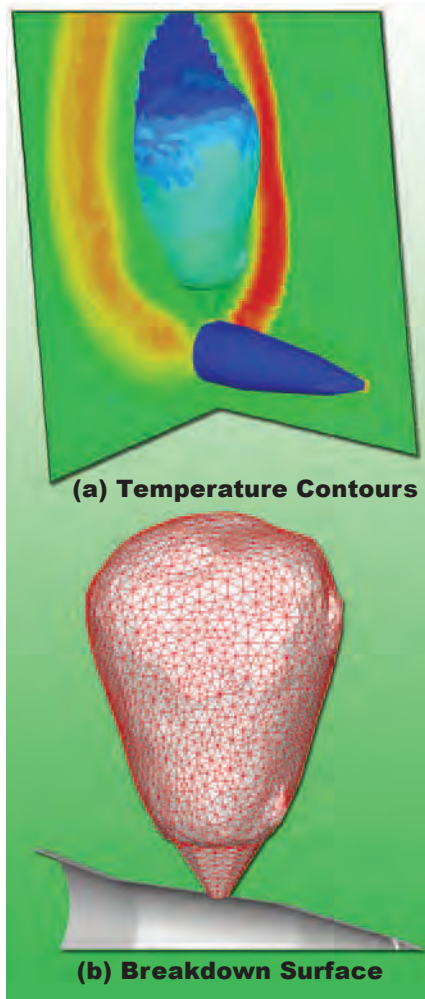


Figure 7. Hybrid solution for missile divert jet problem.

Through the deformation process, the initial seed surface grows until it is constrained by the image (Figures 2c through 2e). Comparison of the final surface to the image constraint (Figure 2f) shows that the fit is quite good.

The capabilities of the AEGIS Toolkit are further demonstrated by generating an interface surface for a more complex multi-nozzle plume (Figure 3). AEGIS integrates the Gnu Triangulated Surface (GTS) toolkit for computational geometry, and GTS is available as freeware.

Two GDM surfaces are seeded and grown as seen in Figure 4. The near-field surface is then processed through the GTS Boolean operator to stitch it to the nozzles (Figure 5). The resulting near-field surface is then merged to the far-field surface (Figure 6) to produce the final interface surface ready for interpolation and conversion to an appropriate DSMC boundary format.

HYBRID SOLUTIONS FOR STEADY FLOWS

A variety of flowfields have been analyzed using the hybrid methodology developed (with coupling making use of breakdown surfaces constructed using AEGIS)

as described in the previous section. The details of such flowfield solutions are available in a number of papers presented at American Institute of Aeronautics and Astronautics (AIAA) meetings over the past several years.^{3,5} A missile divert jet calculation is shown in Figure 7a with the breakdown surface used for continuum/DSMC coupling shown in Figure 7b.

A hybrid solution for an Apollo-like vehicle during re-entry is shown in Figures 8a and 8b. The continuum region is solved using the CRAFT CFD[®] code while the rarefied wake region is solved using PDAC. As can be seen, there is very good continuity across the breakdown surface interface.

Results for an extension of the Apollo re-entry hybrid simulation to three-dimensions are shown in Figures 9a and 9b. Here, three-dimensional effects occur due to angle of attack, as well from the Reaction Control System (RCS) thrusters. For this case, the embedded rarefied wake region, for which a DAC simulation is needed, contains an embedded highly dense RCS thruster core flow.

Continued Next Page...




Figure 6. GTS Boolean operations on GDM generated surfaces.

The versatility of AEGIS allows for the generation of both the outer interface as well as the inner interface.

UNSTEADY HYBRID METHODOLOGY

Extending hybrid methodology to unsteady flows, where the breakdown surface itself may vary with time and where coupling conditions are also transient, has been a challenge. The first step has entailed development of an unsteady version of PDAC that incorporates ensemble averaging of the flowfield properties at discrete time intervals.

Unsteady boundary conditions have also been added to PDAC, which also permits interpolation of a time-varying source and/or wall boundary conditions in a more general manner in order to couple to an unsteady continuum code.

The boundary condition input was modified to permit time variation of each of the source properties on all surfaces. Figure 10 schematizes unsteady coupling methodology on a time-invariant surface with time-varying conditions. A variety of fundamental cases illustrating application of this methodology to high altitude transient jet

expanding domain problems have been examined.

PARTICLE EXTENSIONS

Lagrangian methodology used in our continuum codes² has also been implemented in the PDAC DSMC code. Figure 11 shows an application of methodology for a hybrid CFD/DSMC missile plume simulation.

In a one-way coupled Lagrangian approach, the particle paths and temperatures are determined by using calibrated drag and heat transfer correlations to account for momentum and energy transfer from the gas to the particulates. PDAC uses identical correlations to those used in the continuum methodologies but also uses the same enthalpy-temperature curve fits and emissivity tables as the continuum codes. This allows the particle phase and radiative property changes to occur seamlessly across the coupling boundary.

For large particle mass loadings, significant energy and momentum may be transferred from the solid particulates back to the gas during the plume expansion, and treatment of these problems requires a fully-coupled approach.




Figure 10. Hybrid transient coupling.

Continued Page 18




Figure 8. Apollo-like vehicle simulation.




Figure 9. Apollo-like vehicle re-entry simulation at angle of attack with dual thrusters.

Using ParaView for Remote Visualization on IBM AIX

Sean Ziegeler, HPCMP PET Enabling Technologies On-Site, Mississippi State University

Introduction

ParaView is an open-source application for viewing and analyzing a wide variety of data sets. The major capabilities of ParaView include:

- Support for two-dimensional (2D), three-dimensional (3D), and time-varying data sets.
- Support for structured and unstructured grids.
- Available for many platforms including Windows, Linux, and Mac.
- A client-only mode for running on a local desktop or laptop system.
- A client-server mode for connecting a local system to a server on a remote system.
- A parallel server mode for large data sets.

ParaView has many specific features that are beyond the scope of this article. For more information visit <http://www.paraview.org/> for background information and the

Engineer Research and Development Center (ERDC) Major Shared Resource Center (MSRC) Data Analysis and Assessment Center (DAAC) wiki page on ParaView (<https://visualization.hpc.mil/wiki/ParaView>) for an introduction and tutorials on ParaView in general.

Prerequisites

For this article, we will be using ParaView 3.2.1. If the user does not have this version of ParaView, pre-compiled versions can be obtained for Windows, Linux, and Mac at <http://www.paraview.org/New/download.html>. Users with other platforms will have to compile the software themselves.

Instructions for compiling ParaView can be obtained at http://paraview.org/Wiki/ParaView:Build_And_Install. Fortunately for users of the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC) ParaView is already installed on the

BABBAGE system. Executables for the serial version of ParaView are located in `/site/unsupported/paraview/serial/bin/`. Executables for the parallel Message Passing Interface (MPI) version (currently experimental) of ParaView are located in `/site/unsupported/paraview/mmpi/bin/`.

Client-Server Mode ParaView

ParaView's client-server mode operates, for the most part, transparently to the user. Except for some initial configuration (which is described elsewhere in this article), using ParaView in this way will be just like using it on the user's local system (i.e., in client-only mode).

This article assumes that the user is a beginner, but the user is nonetheless encouraged to become familiar with the basic operation of ParaView in client-only mode, which, in fact, runs a default server on any user's local system (Figure 1).

In client-server mode, the user starts the ParaView server program on a remote system (e.g., BABBAGE, in this case). The user still runs the client on the local system, but that client now connects to the remote server process, as shown in Figure 2.

Secure Shell (SSH) Port Forwarding

In cases where the remote system is behind a firewall, is an inaccessible node on a cluster, or when network traffic should be encrypted, Secure Shell (SSH) port forwarding must be used. This is the case with BABBAGE's compute nodes. This means that the ParaView client connects to SSH, which "tunnels" the traffic through its connection and




Figure 1. (Top) ParaView in client-only mode, running a default server. Figure 2. (Bottom) ParaView in client-server mode between the local system and a remote one.

Continued Next Page...

Code 1. Message indicating job pending.

JOBID	USER	STAT	QUEUE	FROM_HOST	EXEC_HOST	JOB_NAME
235261	seanzig	PEND	standard	b6n1		ParaView

Code 2. Message indicating job is running.

JOBID	USER	STAT	QUEUE	FROM_HOST	EXEC_HOST	JOB_NAME
235261	seanzig	RUN	standard	b6n1	b6n3	ParaView

then forwards it to the ParaView server on the remote system as shown in Figure 3.

Starting the ParaView Server

Running the ParaView server on BABBAGE is just like running any other program. Most importantly, the user must submit a batch job via a script. The following is a Load Sharing Facility (LSF) batch script template that can be used for ParaView.

```
#!/bin/sh
#BSUB -J ParaView
#BSUB -o ParaView%J.out
#BSUB -e ParaView%J.err
#BSUB -P user_account_
number_here
#BSUB -W user_desired_job_
time_here
#BSUB -q user_desired_queue_
here
#BSUB -n 1
/site/unsupported/paraview/
serial/bin/pvserver \
--use-offscreen-rendering --
server-port=23456
```

This script can be tailored to the user's needs in terms of job time, queues, etc. For example, the bigmem queue will be required for large data sets

when running serial ParaView. In this template, the server port is set to 23456, but this can and should be changed by the user to any number greater than 1024. The user must keep this port number in mind for later.

Once the script is submitted, the user must monitor the job to find out when and where it begins running. Using the bjobs command, the user will receive something like Code 1 while the job is still pending (i.e., not running yet).

When the job begins running, the output of the bjobs command will look something like Code 2.

Notice that the status changes from PEND to RUN. Also note that the EXEC_HOST is now specified; it is b6n3 in this case. Keep this host in mind for later. Finally, the bpeek command should be used to verify that the ParaView server is ready to accept connections. If it is ready, it should appear as follows:

Listen on port: 23456

Waiting for client...

Setting up the Secure Shell Port Forward

Using SSH port forwarding is only slightly different than using SSH for a typical command line session.

As usual, the user must have a valid Kerberos ticket for a High Performance Computing Modernization Program (HPCMP) system. There are two programs that can be used—depending upon the user's operating system. In Linux and UNIX, the ssh command is used. In Windows, PuTTY is used. The following command shows how to set up a port forward from the user's local Linux/UNIX system to the respective EXEC_HOST on BABBAGE:

```
ssh -L 23456:b6n3:23456
username@babbage.navo.hpc.mil
```

The port numbers (23456) must match the port number used for the ParaView server. The remote host (b6n3) must match the host provided by the EXEC_HOST column of the bjobs command. After the execution of the above command, the user should appear to be logged into the BABBAGE head node as usual.

The screenshot in Figure 4 shows the port forwarding options in PuTTY for Windows systems. To get to this point and then to fill in the appropriate options, the user must follow these steps:

1. Run PuTTY.
2. Select a Saved Session or create a new session for *babbage.navo.hpc.mil*.



Figure 3. ParaView in client-server mode using SSH port forwarding.

3. Scroll to the bottom of the list on the left of the PuTTY Configuration screen.
4. Click on the Tunnels menu (this should look like Figure 4).
5. Enter the source port: 23456 (match the port number used for the ParaView server).
6. Enter the destination: b6n3:23456 (match the port and EXEC_HOST column of bjobs).
7. Select the Local option (if it is not already selected).
8. Click the Add button.
9. The new forward option should now appear in the list of forwarded ports.
10. Optionally, return to Session menu and save the session.
11. Click the Open button.

After completing the above steps in PuTTY, the user should be connected

to BABBAGE as usual.

Starting the ParaView Client

The final step is to run the ParaView client on the user's local system. To establish a client-server connection based on the above, the user will follow these steps:

1. Run the ParaView client (usually paraview or paraview.exe or by clicking the icon).
2. Click on the File, Connect menu option or the Connect button on the button bar.
3. Click the Add Server button.
4. Give the session a name (e.g., Local Forward).
5. Make sure the Server Type is Client/Server.
6. Make sure the Host is localhost.
7. Set the Port to 23456 (match the port used above).
8. Press the Configure button.

9. Set the Startup Type to Manual.
10. Click Save.
11. Select the new session (it will have the name you gave it in step 4).
12. Press Connect.

Steps 3-10 will only need to be done the first time. ParaView will remember the user's sessions, and from now on, the user need only select that session and press Connect.

If the blank ParaView window reappears without any error messages and something like `cs://localhost:23456` appears in the Pipeline Browser to the left, then the ParaView client on the user's local system is now connected to the server on BABBAGE.

A Simple Test

It is recommended the user test the client-server connection between the local system and BABBAGE for sufficient bandwidth and proper

Continued Page 20

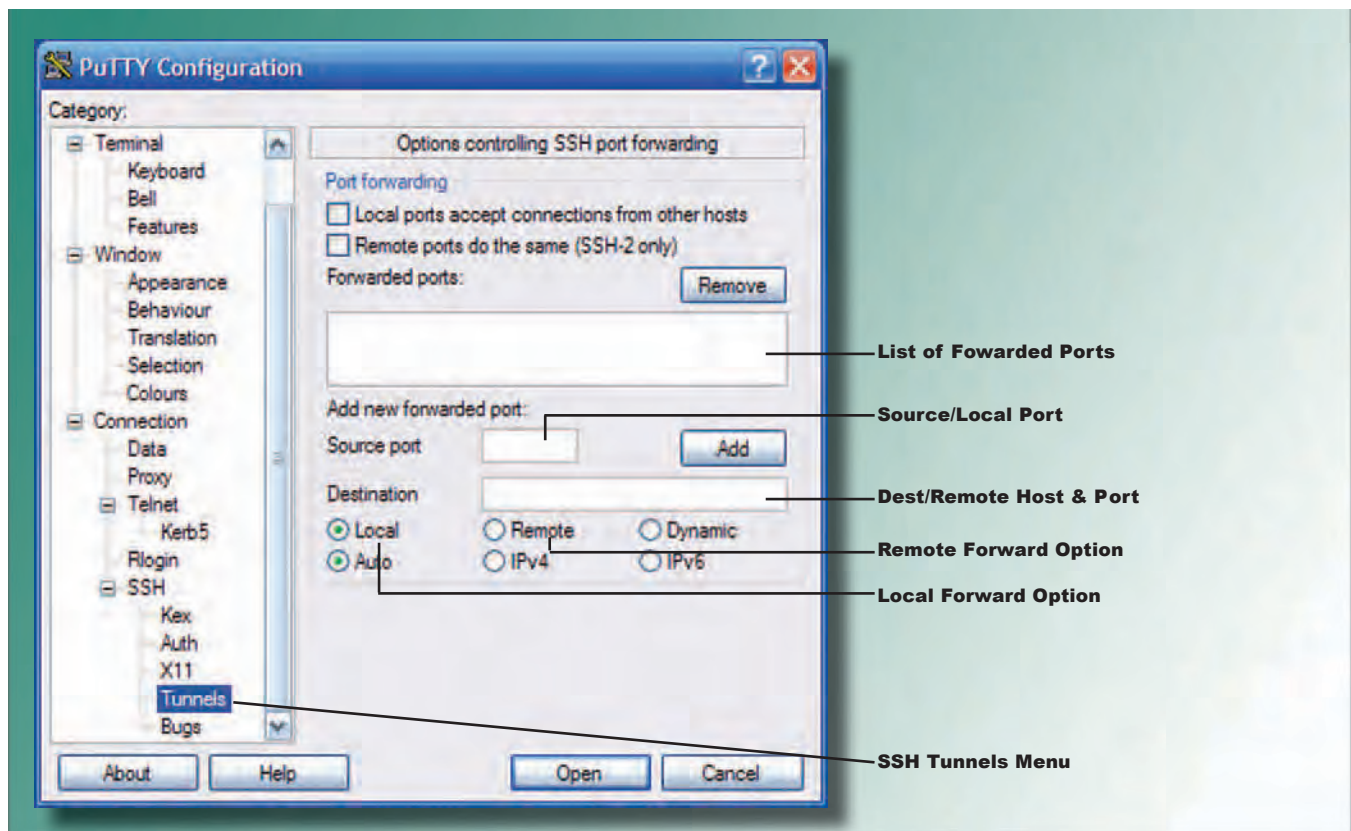


Figure 4. A snapshot of PuTTY with the port forwarding options.

Global and Meso Forecasts of Tr

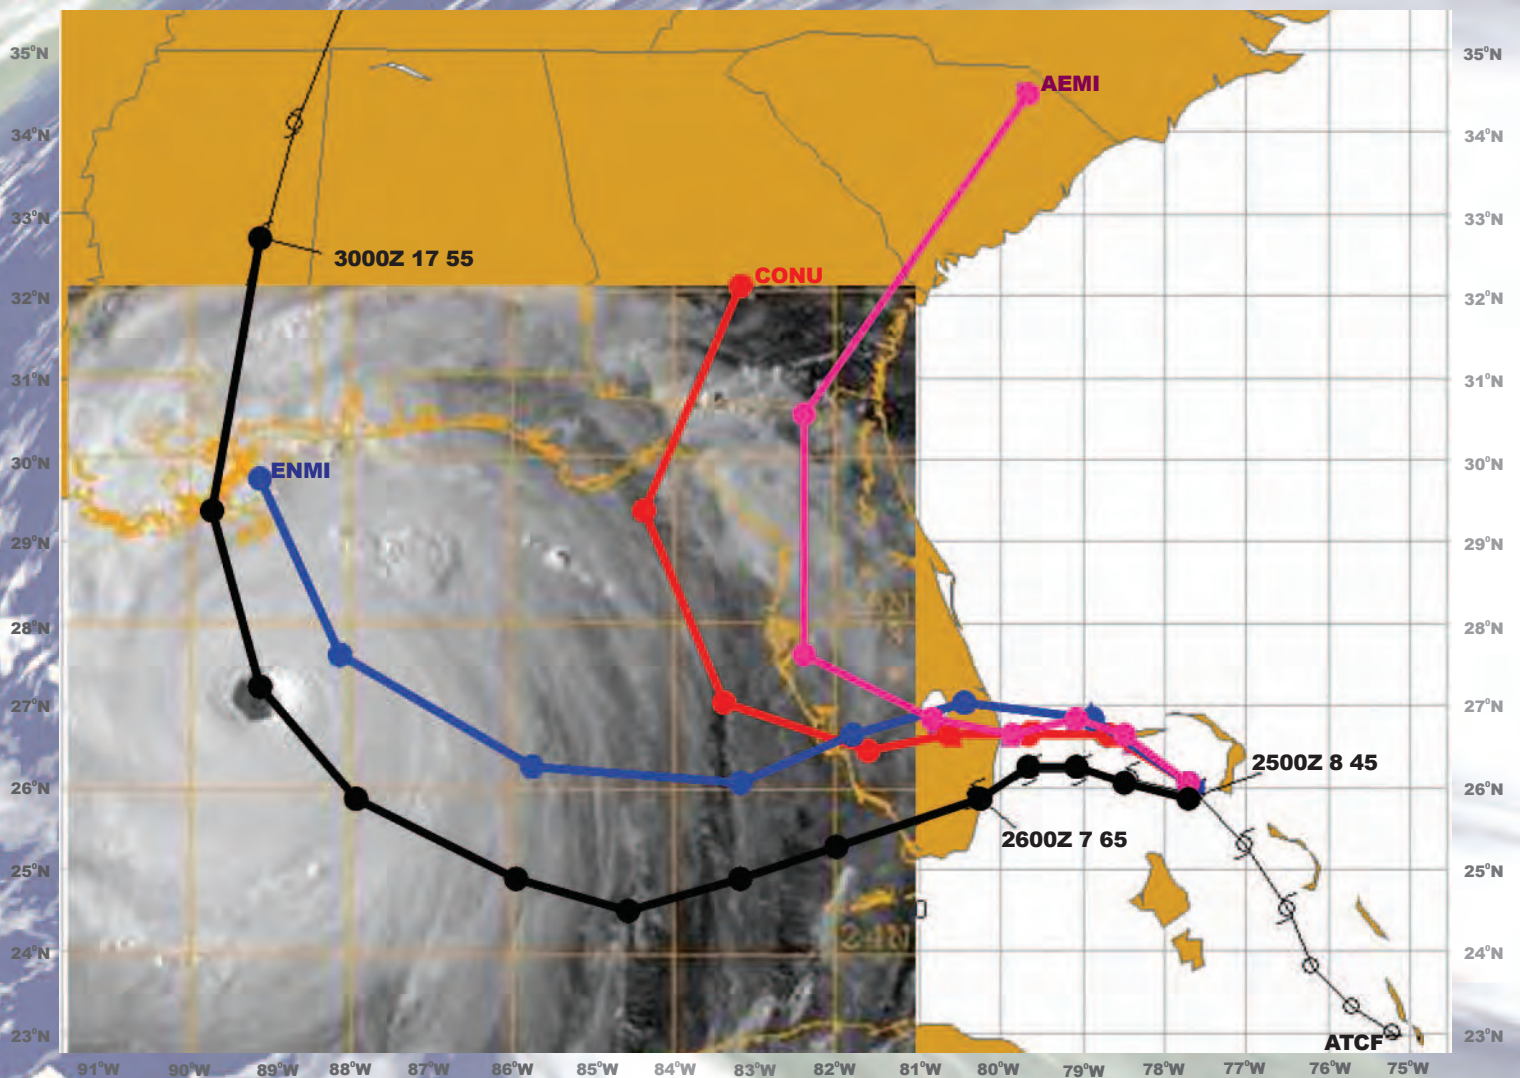


Figure 1. The observed track of Hurricane Katrina (black filled symbols), and the ensemble transform stochastic convection ensemble-mean forecast (ENMI, blue track), multi-model consensus forecast (CONU, red track), and the National Centers for Environmental Prediction Global ensemble-mean forecast (AEMI, pink track), from 25 August 2005. Tracks are superimposed on a visible satellite image of Katrina taken from National Oceanic and Atmospheric Administration Geostationary Operational Environmental Satellite 12 at 2245 UTC on 28 August 2005.

mesoscale Ensemble Tropical Cyclones

C. A. REYNOLDS, J. D. DOYLE, J. GOERSS,
C. AMERCAULT, C. H. BISHOP, T. HOLT AND J. G. MCLAY,
NAVAL RESEARCH LABORATORY, MONTEREY, CA

The accurate prediction of tropical cyclones¹ is of utmost importance to both civilian and Department of Defense (DoD) interests. While research is ongoing to improve deterministic atmospheric forecasts through advancements to the forecast model and more accurate estimates of the initial state, there is also interest in obtaining probabilistic information derived from ensemble forecasts. Ensembles of forecasts from equally plausible initial states and model formulations offer a computationally feasible way of addressing inevitable forecast uncertainties. Ensembles offer improved forecasts through ensemble mean quantities and quantitative estimates of forecast error. Although the concept of ensemble modeling is relatively simple, the performance of an ensemble forecast system is very sensitive to the basic ensemble architecture.

Continued Next Page...

1. Strong tropical cyclones are known as hurricanes in the Atlantic and eastern Pacific and as typhoons in the western Pacific.

At the Naval Research Laboratory, we are designing new ensemble methods for both the global and mesoscale atmospheric forecast systems. Because of the high computational demands associated with ensemble development and verification (especially when one is interested in severe or rare events), exceptional computational resources are necessary to perform this research.

The development of new operational global and mesoscale atmospheric ensemble systems has been made possible through the High Performance Computing (HPC) Challenge Project entitled Multi-scale Predictability of High-impact Weather in the Battlespace Environment. The computations have been performed primarily on the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC) IBM Power 4+ (KRAKEN) and IBM Power 5+ (BABBAGE) systems.

The global atmospheric ensemble forecasts are produced using the Navy Operational Global Atmospheric Prediction System (NOGAPS) atmospheric model, a global spectral model with a hybrid-sigma coordinate in the vertical. The prognostic variables are vorticity, divergence, virtual potential temperature, specific humidity and terrain pressure.

The parameterizations of sub-grid-scale physical processes include deep convective precipitation, shallow cumulus mixing, convective and stratiform clouds, boundary layer vertical mixing, surface flux parameterization, gravity wave drag, and radiation parameterization. The ensemble is run at a horizontal resolution of approximately 1 degree. In this study, each ensemble contains 32 perturbed forecasts in addition to one control forecast.

For the results described here, the initial perturbations are created using a new technique, the Ensemble Transform (ET), which combines information from the operational data

assimilation system with short-term forecasts to produce perturbations to the initial state that are both growing and consistent with the estimate of analysis error produced by the data assimilation system. Five-day ensemble forecasts are performed twice daily for the period of July through October 2005. These computations required approximately 50,000 Central Processing Unit (CPU) hours on KRAKEN and BABBAGE.

To account for model error, stochastic perturbations are added to the tendencies produced from the parameterization of deep convection, which represent uncertainty in the impact of the sub-grid-scale convective clouds on the resolved scales. These

ensembles are compared to experiments in which only the initial state was perturbed (i.e., no stochastic convection). The ensemble experiments exhibit clear evidence of improved ensemble performance under a variety of measures when they include stochastic convection. In fact, the NOGAPS ensemble mean tropical cyclone track forecasts are competitive with high-resolution multi-model forecasts (an ensemble composed of high-resolution forecasts from different operational centers) at longer lead times.

The multi-model forecast ensemble is traditionally considered one of the leading track guidance tools for forecasters. As an example, we consider the case of Hurricane




Figure 2. The average forecast error in nautical miles for the 2005 Atlantic Hurricane season for the high-resolution NOGAPS deterministic forecast (NGPI, blue), the ensemble transform stochastic convection ensemble-mean forecast (ENMI, red), the multi-model consensus forecast (CONU, green), and the official forecast (OFCL, pink). The number of forecasts included range from 358 at 24 hours to 133 at 120 hours.

Katrina, which was the costliest and one of the five deadliest hurricanes to ever strike the United States. Figure 1 shows the observed track of Hurricane Katrina, as well as 5-day forecasts from the NOGAPS ensemble mean (blue track), multi-model consensus (red track), and the National Centers for Environmental Prediction Global ensemble-mean (pink track), from 25 August 2005.

Clearly the NOGAPS ensemble forecast is significantly better than the other forecasts at long lead times in this case. Examination of the entire 2005 season shows that, on average, the new NOGAPS ET ensemble with stochastic convection performs comparably or better than multi-model consensus forecasts and the official forecast track after about 3 days (Figure 2 shows average hurricane track forecast errors for the Atlantic basin).

The mesoscale ensembles are produced using the atmospheric module of the Coupled Ocean Atmosphere Mesoscale Prediction

System (COAMPS®), which is a nonhydrostatic model capable of accurate prediction at high resolution and makes use of a terrain following vertical coordinate.

The model is a finite-difference approximation to the fully compressible, nonhydrostatic equations. Physical parameterizations are used to represent surface fluxes, boundary layer, radiation, and moist processes including microphysical quantities including cloud water, cloud ice, snow, rain, graupel/hail, and water vapor.

The COAMPS ensemble system makes use of the ET method to provide initial condition perturbations, and the NOGAPS ET ensemble provides dynamically-consistent perturbations through the lateral boundaries that represent the synoptic-scale uncertainties.

The COAMPS ensembles are run with a 45 kilometer (km) resolution with 28 members with a forecast length of 48 hours. Model uncertainty is introduced through variations in

parameters within the various physical parameterizations based on typical uncertainty ranges, focusing on moist processes and the boundary layer.

Figure 3 shows the 5840-meter geopotential height isopleth for the 48-hour COAMPS mesoscale ensembles run for Hurricane Dennis, initialized at 00 UTC (Coordinated Universal Time) 9 July 2005. There is some ensemble spread in the ensemble with initial perturbations only (left panel), particularly in the vicinity of Hurricane Dennis. The COAMPS ensemble with both initial perturbations and parameter variations, right panel, shows a significant increase in ensemble spread, particularly in the tropics, that translates to spread in the track and intensity forecast for Dennis.

The recently developed adjoint and tangent linear model system for COAMPS provide a new opportunity to explore mesoscale predictability of tropical cyclones. The adjoint model provides a sophisticated method to quantitatively compute the forecast

Continued Next Page...

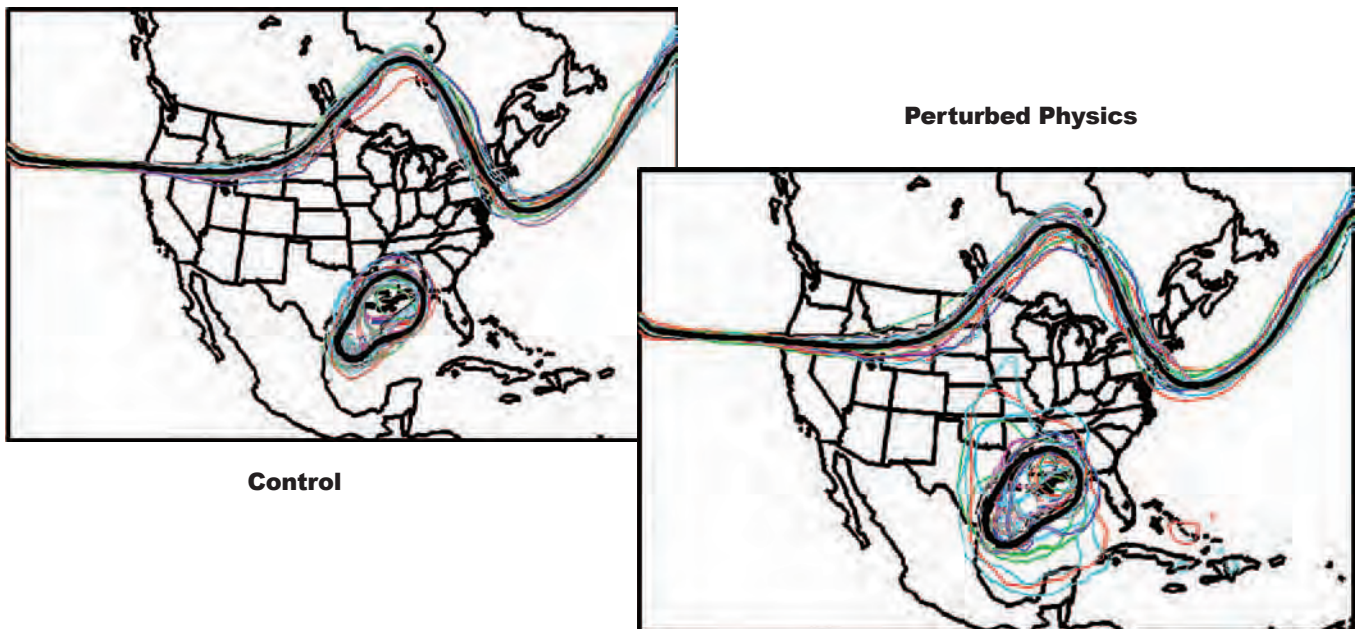


Figure 3. Ensemble isoheights (different color for each ensemble member) and ensemble mean (black) 48-hour COAMPS forecasts for Hurricane Dennis, from 00 UTC 9 July 2005 for the 5840-meter isoheight. The ensemble with perturbed physics (right) shows considerably more ensemble variance than the control ensemble without perturbed physics (left).

sensitivity. A unique capability of the COAMPS adjoint modeling system is that an accurate representation of the moist processes is achieved through an explicit adjoint of the microphysical parameterization.

Figure 4 shows an example of the application of the adjoint system for Hurricane Katrina. The sensitivity of forecast energy in a box surrounding Katrina at 00 UTC 29 August 2005 to the previous 24-hour model vorticity at 2.5 km is shown in the elevated horizontal surface. The sensitivity pattern indicates a complex interwoven banded structure,

suggestive of vortex instabilities present within Katrina that may limit the intrinsic predictability.

The sensitivity results indicate a subtle interplay between moist convection, vortex dynamics, and air-sea interaction instability, all of which conspire to limit predictability of tropical cyclones such as Katrina on the smaller scales.

Prediction of tropical cyclone track and intensity remains one of the greatest challenges in meteorology today. The results of this research highlight the promise of ensemble-based approaches for tropical

cyclone prediction using global and mesoscale models.

This research on ensemble design, made possible through the DoD HPC challenge program, has already led to significant improvements in the Navy's global atmospheric ensemble prediction system (the ET was transitioned to operations in March 2008). In addition, this research will lead to new capabilities in the form of mesoscale ensemble forecasts, providing the Navy with probabilistic forecasts of strategically important weather parameters on the tactical scale for the first time.

Acknowledgments

The support of the Naval Research Laboratory and the Office of Naval Research under program elements 0601153N and 0602435N is gratefully acknowledged. Computational resources were provided through the DoD HPC Challenge Project "Multi-scale Predictability of High-impact Weather in the Battlespace Environment."




Figure 4. Three dimensional depiction of the forecast sensitivity based on a COAMPS forecast of Hurricane Katrina using the adjoint and tangent linear model system. The sea-level pressure (white contours) and 10-m wind speed (color) are shown at the surface. The sensitivity of the energy in a box surrounding Katrina at 00 UTC 29 August 2005 to the previous 24-hour model vorticity at 2.5 km is shown in color elevated above the surface. The three dimensional surface corresponding to the equivalent potential temperature of 340 K and shaded by wind speed is also displayed.

EINSTEIN and DAVINCI Come to the MSRC

Christine Cuicchi, Computational Science and Applications Lead, NAVO MSRC

The Technology Insertion–Fiscal Year (FY) 2008 (TI-08) procurement process provides the Naval Oceanographic Office Major Shared Resource Center (NAVO MSRC) with a departure from recent years while continuing the line of well-established, highly available IBM Power systems that have dominated the MSRC's infrastructure for nearly eight years.

The MSRC is preparing for the arrival of an 80 Teraflop IBM P6 cluster and a 117 Teraflop CRAY XT5 cluster in June and August, respectively.

The addition of these powerful systems will quadruple the NAVO MSRC computational capacity by the close of 2008: EINSTEIN (XT5) and DAVINCI (P6) will be joining BABBAGE and PASCAL in providing an aggregate 233 Teraflops of computational capability to our users.

Both systems will also feature a small subset of big memory nodes, and both will employ the PBS Professional™ scheduler, a departure from the current workload management system on our existing systems: Load Level Facility (LSF).

As our new systems are installed and tested, we'll also be retiring our P4+ systems KRAKEN and ROMULUS starting on 1 October 2008. As the decommission date nears, we will provide users with more information regarding queue draining and movement of data from the P4+ systems.

We expect that EINSTEIN and DAVINCI will be available for production use in late 2008. For updates on our new systems' availability, please visit <http://www.navo.hpc.mil>.

	IBM P6 DAVINCI	CRAY XT5 EINSTEIN
Processor Type	IBM Power6	AMD Quad-core Opteron
Processor Speed	4.7 GHz	2.3 GHz
Gflops/core	18.8	9.2
Peak Performance	80 Teraflops	117 Teraflops
Total Hard Disk Capacity	437 Terabytes	518 Terabytes
Number of Compute Nodes/Cores	133 Nodes 4,256 Cores	1,592 Nodes 12,736 Cores
Cores per Node	32	8
Number of Big Memory Nodes	2	8
Memory Available to User, Standard Memory Nodes	1.7 Gb per core 54.4 Gb per node	1.9 Gb per core 15.8 Gb per node
Memory Available to User, Big Memory Node	7.7 Gb per core 246.4 Gb per node	3.9 Gb per core 31.8 Gb per node
Processor to Memory Bandwidth	512 Gb/s (peak per node)	10.67 Gb/s/Quad-core 2.67 Gb/s/core
Interconnect	Infiniband	Seastar2
File System	GPFS	Lustre
Operating System	AIX	Linux
Total System Size with Service Nodes	4,768 cores	12,872 cores

Breakdown of DAVINCI and EINSTEIN systems.

Hypersonic Continuum...continued

Fully-coupled methodologies have been implemented in the continuum codes,² and fully-coupled models are currently being developed for implementation in the PDAC DSMC code as well. A simple, fully-coupled heat-bath model has been developed, and results discussed in recent papers show that the model provides the correct thermodynamic behavior in a simple DSMC simulation.

CONCLUSION


A unique hybrid continuum/DSMC methodology has been developed to enable accurate predictions of plume flowfields and flow over re-entry vehicles at altitudes in which rarefaction effects are important but significant parts of the domain can be described by continuum methods.

A key feature of this methodology is the AEGIS toolkit, which permits the construction of breakdown surfaces for complex flows in an automated manner. These complex flowfields

may include multiple continuum regions, and thus multiple interface surfaces. Extensions to the hybrid methodology for transient problems and for analyzing particulate laden plumes/jets in rarefied zones are also included.

Future work will continue to address many of the computational complexities in modeling these types of flows and the additional challenges posed by transient flow phenomena, which include geometric changes.

Figure 11. Hybrid two-phase PDAC solution for 120 km plume.



Acknowledgements

This work was supported by the U.S. Army Aviation and Missile Research, Development, and Engineering Center (AMRDEC) under the direction of Dr. Billy Walker. High Performance Computing Center (HPCC) resources were made available under a program entitled "Altitude Effects – Fluid Flow Transition and Continuum Breakdown" and varied platforms were used during the course of this program.

References

1. Bird, G.A., "Molecular Gas dynamics and the Direct Simulation of Gas Flows," Oxford Science Publications, Clarendon Press, Oxford, 1994.
2. Papp, J.L., B.J. York, N. Sinha, and S.M. Dash, "Progress in Modeling Particle Turbulent Dispersion in High-Speed Jets/Plumes," AIAA-2003-1284, 41st AIAA Aerospace Conference and Exhibit, Reno, NV, 6-9 January 2003.
3. VanGilder, D., C.C. Chartrand, J.L. Papp, R. Wilmoth, and N. Sinha, "Computational Modeling of Nearfield to Farfield Plume Expansion," Paper AIAA-2007-5704, 43rd AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, Cincinnati, OH, 8 – 11 July 2007.
4. LeBeau, G.J. and F.E. Lumpkin III, "Application Highlights of the DSMC Analysis Code (DAC) Software for Simulating Rarefied Flows," Computer Methods in Applied Mechanics and Engineering 191 (6-7): 595-609, 2001.
5. Papp, J.L., R. Wilmoth, C. Chartrand, D.C. Kenzakowski, and N. Sinha, "Simulation of High-Altitude Plume Flow Fields Using a Hybrid Continuum CFD/DSMC Approach," AIAA-2006-4412, 42nd AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, 9 – 12 July 2006, Sacramento, CA.



The Porthole



Visit of LCDR Mike Cooper and Ms. Jane Bruno, OPNAV N43C, to the NAVO MSRC.



Bomb Squad from Biloxi, MS.



The Basic Oceanography Accession Training (BOAT) class students visit the NAVO MSRC.



IBM representatives prepare for installation of TI08 POWER6 system.



NASA representatives visit the NCCIPS Data Center.



Faculty from Mississippi State University.

ParaView...continued

configuration. The following steps provide a straightforward test:

1. Click Sources Wavelet (to create a “fake” 3D data set).
2. In the panel on the left, leave all options as default and click the Apply button.
3. An outline of a cube should appear in the main window.
4. On the second drop-down box from the top-left, change it from Outline to Surface.
5. The cube should become solid.
6. On the first drop-down box from the top-left, change it from Solid Color to RTData.
7. Click and drag with the left mouse button in the main window to rotate the cube.

The cube should look something like Figure 5.

At this point the user can apply other operations to the fake data, e.g., contours (isosurfaces) or cut planes, to determine if the performance will be sufficient.

Other Options

If the performance is not quite acceptable, some client-server options

can be fine-tuned to trade off quality for speed. Such options are available in the Edit Settings menu under the Render View: Server screen.

The “Remote Rendering Parameters” determine when ParaView switches from local rendering to remote rendering. The “Client/Server Parameters” determine how much compression is used for sending interactive images between the client and server. Higher settings will result in greater compression but lower image quality.

Running a Parallel Server

In cases where the data sets are very large, only a parallel, i.e., Message Passing Interface (MPI), server will suffice to maintain interactive rendering rates. Some aspects of the parallel server are experimental at this time, but in some cases it is still usable. The following script template shows how to run a parallel ParaView server job:

```
#!/bin/sh
```

```
#BSUB -J ParaView
```

```
#BSUB -o ParaView%J.out
```

```
#BSUB -e ParaView%J.err
```

```
#BSUB -P user_account_  
number_here
```

```
#BSUB -W user_desired_job_  
time_here
```

```
#BSUB -q user_desired_queue_  
here
```

```
#BSUB -n user_desired_num_  
processors
```

```
/site/unsupported/paraview/mpi/  
bin/pvserver \
```

```
--use-offscreen-rendering --  
server-port=23456
```

The two primary differences are (1) the user may now select the desired number of processors to use and (2) the path to the pvserver executable now contains “mpi.”

For now, when using a parallel server, remote rendering must be disabled; otherwise, the user will see a black screen. In the Edit, Settings menu, under the Render View: Server screen, the “Remote Render Threshold” check box must be turned off.

Conclusion

This specific procedure for using the client-server capabilities of ParaView was written specifically for the BABBAGE system. However, it can be adapted to other systems assuming ParaView is installed. Other platforms, operating systems, and batch queuing programs have equivalent commands to gather the information needed to establish the remote connections.

ParaView is a versatile visualization and analysis program. As demonstrated in this article, the remote visualization capability is especially powerful. By following the instructions and tips provided in this article, users should be able to harness this power.

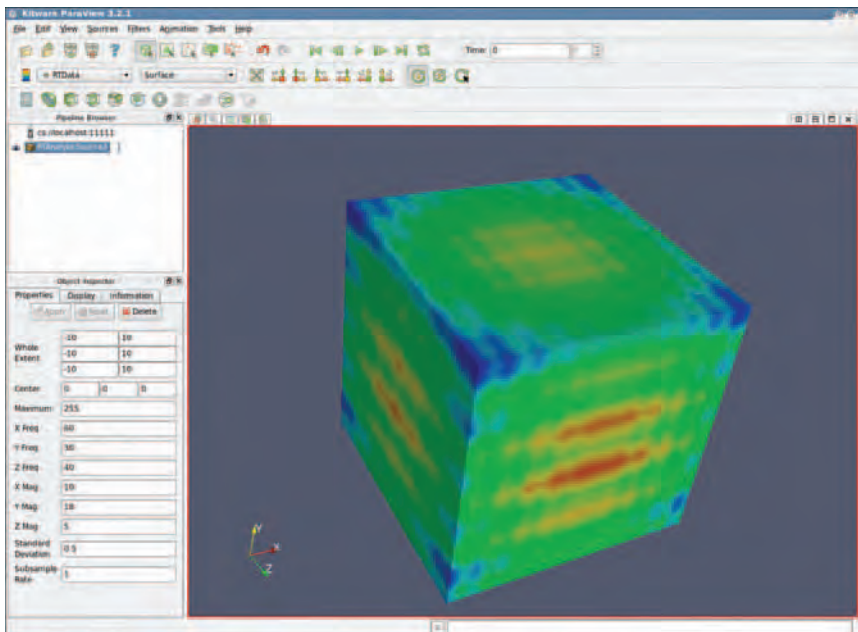


Figure 5. A snapshot of ParaView with the “fake” data.

Top 5 Frequently Asked Questions about \$HOME and \$WORKDIR on the NAVO MSRC High Performance Computing Systems

Sheila Carbonette, NAVO MSRC User Support

This issue of Tips 'n Tricks is designed to provide answers to some of our users' most frequent questions on disk storage and usage with \$HOME and \$WORKDIR on the Naval Oceanographic Major Shared Resource Center (NAVO MSRC) High Performance Computing (HPC) systems.

Q. How do I check the disk usage for my home and work directories on KRAKEN and BABBAGE?

A. The `/site/bin/quota` command can be used on the IBM systems to check default quotas and current disk usage. An example is shown in Code 1.

The current default user quotas for \$HOME and \$WORKDIR on KRAKEN and BABBAGE are shown below:

	\$HOME	\$WORKDIR
KRAKEN	500MB	5TB
BABBAGE	1GB	10TB

As can be seen, \$HOME is set to a much smaller amount than \$WORKDIR. It is intended to be used to store routine scripts and/or code. \$WORKDIR, on the other hand, is much larger and is intended to be used for temporary work space for input/output data files from running jobs.

Q. How do I request an increase in my default disk quota?

A. Requests for quota increases can be submitted to the Help Desk for consideration. The following information should be included in the request to be evaluated by NAVO MSRC management:

- ✧ How much additional space is needed.
- ✧ How long the increase is needed.
- ✧ Brief description of why the increase is needed.

Q. How do I get a file restored that has been accidentally deleted from my home directory on KRAKEN?

A. The \$HOME and \$WORKDIR file-systems are not backed up. If a file or directory is deleted, or if there is a major disk failure, these files or directories are lost. Users should back up working files in \$HOME and \$WORKDIR to the archive servers, JULES or VINCENT, or to some other permanent storage location. Please contact the Help Desk if assistance is needed in creating backups of these directories.

An example of an Load Sharing Facility (LSF) transfer job to copy files from the archive server, JULES, to \$WORKDIR, `(/scr/shecar)` on BABBAGE is:

```
#!/bin/csh
#BSUB -J transjob      # Name of the Job.
#BSUB -o %J.out        # std output file name
                        # (%J is the Job ID)
```

b6n1% /site/bin/quota

Filesystem type	KB	Block Limits			File Limits					remarks	
		quota	limit	in_doubt	grace	files	quota	limit	in_doubt		grace
gpfs_hm USR	16	1000192	1048576	0	none	1	0	0	0	none	
gpfs_wk USR	192	10000000000	10737418240	0	none	3	0	0	0	none	

Code 1. `/site/bin/quota` command used to check default quotas and current disk usage.


```

#BSUB -e %J.err      # std error file name
#BSUB -W 2:00        # Wall clock time of 2
                      # hours
#BSUB -n 1           # Specify the number of
                      # nodes
#BSUB -P NAVOSLMA    # Project identifier
#BSUB -q transfer     # Queue name
#
# Copy required workfiles from JULES to
$WORKDIR on BABBAGE
#
cd $WORKDIR
pwd
rcp jules:babbageinput.tar .
ls -ltra
# End of script -----

```

while below is an example of an LSF transfer job to copy output files from \$WORKDIR to the archive server, JULES:

```

#!/bin/csh
#BSUB -J transjob     # Name of the Job.
#BSUB -o %J.out       # std output file name (%J
                      # is the Job ID)
#BSUB -e %J.err       # std error file name
#BSUB -W 2:00         # Wall clock time of 2
                      # hours
#BSUB -n 1           # Specify the number of
                      # nodes
#BSUB -P NAVOSLMA    # Project identifier
#BSUB -q transfer     # Queue name
#
# Copy output files from $WORKDIR on
BABBAGE to a subdirectory on JULES
#
cd $WORKDIR
pwd
rcp *.out jules:/u/home/shecar/transfer_run1
ls -ltra
# End of script -----

```

Q. What is the best way to transfer files from \$HOME and/or \$WORKDIR to my home directory on the archive servers?

A. There are a number of ways to transfer data from the compute servers to the local archive servers or to other remote non-NAVO systems:

Transfers to JULES or VINCENT:

- ✧ * Non-Kerberized rsh/rcp from the login nodes (interactive)
- ✧ * Kerberized rsh/rcp/ftp from the login nodes
- ✧ * rsh/rcp or krsh/krpc/kftp from a LSF transfer queue job

Note: Transfer queue jobs run on the login nodes, so both rcp and krpc can be used. However, users can only use non-kerberized rsh/rcp from the compute node internal network (accessible only via LSF batch jobs).

Transfers to other remote non-NAVO systems:

- ✧ * krpc/krsh/kftp from login nodes (interactive)
- ✧ * krsh/krpc/kftp from an LSF transfer queue job

Note: There is no network routing from the compute node internal network on the IBM systems to any NAVO MSRC systems except the archive servers. There is no network routing to any off-site systems from the IBM system's internal network.

Q. How long can I store files in \$WORKDIR on KRAKEN and BABBAGE?

A. \$WORKDIR is defined as a local temporary work space where files and directories are subject to deletion, depending on their modification date/time. At the NAVO MSRC, files with a three to five days old modification date/time are subject to be deleted.

The disk usage of the /scr file-system is evaluated four times a day using the following capacity percentages with Retention Periods (RPs) of three to five days:

- ✧ /scr is >= 85% full --> we will make a pass with the scrubber using a RP of 3 days.
- ✧ /scr is >= 50% but < 85% full --> we will make a pass with the scrubber using a RP of 4 days.
- ✧ /scr is < 50% full --> we will make a pass with the scrubber using a RP of 5 days.

It is the user's responsibility to transfer needed files to the archive servers, JULES or VINCENT, or to some other permanent file storage location.

Coming Events

SNPD08

9th IEEE/ACIS International Conference on Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing * 6-8 August 2008 * Phuket, Thailand

MSC08

IEEE Multi-conference on Systems and Control * 3-5 September 2008 * San Antonio, TX * <http://conferenze.dei.polimi.it/msc08/index.html>

IST08

2008 IEEE International Workshop on Imaging Systems and Techniques * 10-12 September 2008 * Chania, Island of Crete, Greece * <http://ist.ieee-ims.org/index.php>

GridNets08

Second International Conference on Networks for Grid Applications * 8-10 October 2008 * Beijing, China

ICCD08

* <http://gridnets.org/> International Conference on Computer Design * 12-15 October 2008 * Lake Tahoe, UT * <http://www.iccd-conference.org/>

VIS08

IEEE Visualization 2008 * 19-24 October 2008 * Columbus, OH * <http://vis.computer.org/VisWeek2008/Vis/index.html>

CollaborateCom

4th International Conference on Collaborative Computing * 13-16 Nov 2008 * Orlando, FL * <http://www.icst.org/>

SC08

Supercomputing 2008 * 15-21 November 2008 * Austin, TX * <http://sc08.supercomputing.org/>



Naval Oceanographic Office * Major Shared Resource Center
1002 Balch Boulevard, Stennis Space Center
Mississippi 39522